

The Temporal Logic of the Tower Chief System

Lyman R. Hazelton, Jr.*
Flight Transportation Laboratory
Department of Aeronautics and Astronautics
Massachusetts Institute of Technology
Cambridge, MA

1 Introduction

Human reasoners engage in diverse and powerful forms of cognition. Perhaps the most common mechanism used by humans is induction, which is the usual process of learning. The most celebrated and well understood process of thought is deduction, the method employed in mathematical proof. Humans can also make plans, reason about events and causality in an environment which frequently changes and often provides incomplete information. Humans need not be certain about their facts, and can make assumptions when required information is unavailable.

Classical logic deals almost exclusively with deduction. In chapter 4 of Aristotle's *Metaphysics* [1], there is a section regarding the application of logic to predicting the future. In that treatise, Aristotle describes what has now become known as the axiom of Excluded Middle. In essence, he states that in his logic a proposition can be either TRUE or FALSE, exclusively. No other values for the veracity of a proposition are possible. There is no middle ground. Because the future is indeterminate, one cannot assign absolute knowledge of a forthcoming event. This is a broad statement which has far reaching implications. As pointed out by Bertrand Russell in the early nineteen hundreds, it is not a trifling thing to easily toss aside.

On the other hand, perhaps the interpretation of the axiom of excluded middle has been too broad. Aristotle was seeking truth in a very absolute way. He was quite aware of the works of Zeno and Pythagoras, and to a certain extent his logic was developed from his knowledge of their process of mathematical proof. Propositions in mathematics are *universal* in their temporal extent. One never hears a geometer state that two lines are parallel from two until four this afternoon. The lines are simply parallel or they are not. The geometer's proof makes no reference to time, and, in like manner, neither does Aristotle's logic.

One should not lose track of the fact that logic is primarily a model for analysis of human thought. Its utility as a paradigm for symbolic reasoning and computer programming is secondary to its value as a model of the cognitive process. The restriction to exclude time simplifies the model and makes it much more comprehensible. The restriction also limits the range of problems to which the model may be successfully applied.

Clearly, human reasoners often contemplate the future and create rational plans about it. During the last century, logicians have begun to extend classical logic in order to

*Research supported by NASA contract number NGL 22-009-640.

better model this kind of cognition. *Modal logics* have been invented to describe systems in which propositions may or might be true or false. *Defeasible logics* and *model theory* are used for systems which involve change. And *temporal logics* have been devised to reason about time. The purpose of this paper is to describe the logic used in the reasoning scheme employed in the *Tower Chief* system.

2 Classical Logic

Let us begin with a short review of the fundamental ideas of logic. The elementary operands of logic are propositions. A proposition is a declarative statement such as, “Socrates was a man”, or “There is ice on runway 33.” Logic consists of:

1. A set of abstract operations for combining propositions in ways that preserve the veracity of the resulting compound proposition;
2. A mechanism for generating new propositions from (i) knowledge already known and (ii) general statements called *rules*. This process is called *deduction*.

Symbolically, a proposition which is believed to be true may be represented as p . The belief that the same proposition is not true (false) is represented as $\sim p$. The opposite of a proposition being false is that the proposition is true, or $\sim(\sim p) \equiv p$. A situation in which p and $\sim p$ are claimed to be true is a *contradiction*. In classical logic, contradictions indicate an error in the logical system in which they occur.¹

Propositions may be combined via conjunction (logical AND) or disjunction (logical OR). The symbolic representation of the conjunction of two propositions is:

$$p \wedge q \tag{1}$$

and that of the disjunction² is:

$$p \vee q \tag{2}$$

The laws of commutation and distribution for logical combination operators are respectively:

$$(p \wedge q) \equiv (q \wedge p) \tag{3}$$

$$(p \vee q) \equiv (q \vee p) \tag{4}$$

and

$$(p \wedge (q \wedge r)) \equiv (p \wedge q) \wedge (p \wedge r) \tag{5}$$

$$(p \wedge (q \vee r)) \equiv (p \wedge q) \vee (p \wedge r) \tag{6}$$

$$(p \vee (q \vee r)) \equiv (p \vee q) \vee (p \vee r) \tag{7}$$

$$(p \vee (q \wedge r)) \equiv (p \vee q) \wedge (p \vee r) \tag{8}$$

¹Generally caused by an erroneous *rule*. Read on.

²The use of the symbol \vee for disjunction is from the Latin word *vel*, meaning “inclusive OR”. Unlike English, Latin possesses a separate word for “exclusive OR”, *aut*.

Any combination of two propositions may be treated as a single proposition.

Negation applied to combinations leads to the following relationships, known as DeMorgan's Theorems

$$\sim (p \wedge q) \equiv (\sim p \vee \sim q) \quad (9)$$

$$\sim (p \vee q) \equiv (\sim p \wedge \sim q) \quad (10)$$

2.1 Rules

Deductive inference is accomplished through the use of *rules*.

“A *Rule* is a hypothetical proposition composed of an antecedent and consequent by means of a conditional connective or one expressing reason which signifies that if they, viz. the antecedent and consequent are formed simultaneously, it is impossible that the antecedent be true and the consequent false.”³

Translated into more modern terms, a rule is a conditional statement, consisting of an antecedent (the set of conditions to be met) and a consequent (the set of inferences to be implied if the conditions required in the antecedent are satisfied).

While the laws of logic, such as the commutative law above, and the theorems derivable from them, such as DeMorgan's Theorems, are *domain independent*, rules are based on semantic information. Thus,

$$\text{“If it is raining, then the runways are wet.”} \quad (11)$$

is a rule. Such a rule is written formally as

$$p \supset q \quad (12)$$

where p represents the antecedent (“it is raining”) and q represents the consequent (“the runways are wet”). The hypothetical nature of the statement is embodied in the symbol \supset .

Generally, rules have more than one antecedent, and may have more than one consequent. Rules may combine antecedents purely by conjunction, purely by disjunction, or in combination. As an alternative, a disjunctive rule can be split into several simple or conjunctive rules. While this is less efficient from a notational point of view, it avoids disjunction altogether. This is valuable when creating a computer program to do logic, because the program need only perform conjunction.

There are six forms in which deductive rules of inference may appear. For the purposes of this discussion, the description of two will suffice.

³This definition is from a translation of the fourteenth century logician Pseudo-Scotus (John of Cornubia) which appears in [1].

Modus Ponens This form, properly known as *modus ponendo ponens*, or the method of affirmation leading to affirmation⁴, comes to the conclusion of the consequent if the hypothetical antecedent is declared. For the rule above, the information that it is raining leads to the inference that the runways are, indeed, wet. Formally

$$\begin{array}{l} p \supset q \\ p \\ \hline \therefore q \end{array} \quad (13)$$

Modus Tolens The correct name is *modus tollendo tollens*, or the method of denial leading to denial.⁵ In this form, the negation of the antecedent is concluded if the hypothetical consequent is declared **not** to be true. The formal definition is

$$\begin{array}{l} p \supset q \\ \sim q \\ \hline \therefore \sim p \end{array} \quad (14)$$

A *rule* may be thought of as a generalization that can be applied to a domain of specific situations. In effect, rules are the analogs of algebraic equations in logic. This generalization is accomplished, as in algebra, through the use of *variables*.

Note that the rule stated in (11) is not really correct. Obviously, if the rain is falling in Boise, there is little effect on the runways at LaGuardia. The rule as stated is not precise enough. A more precise statement of the intended meaning of rule (11) is

$$\text{"If it is raining at some airport, then the runways at that airport are wet."} \quad (15)$$

If the airport is represented by the variable x , that it is raining at x by Rx , and that the runways are wet at x by Wx , then the formal statement of rule (15) corresponding to (12) is

$$Rx \supset Wx \quad (16)$$

Further information on topics such as quantification can be found in [2].

3 Defeasible Logics and Truth Maintenance

In a classic paper in 1979, Jon Doyle [3] made a first attempt to create a computer program employing an extended classical logic for use with dynamic domains. Doyle's *Truth Maintenance* is based on the idea that reasoned inferences are supported by evidence. The evidence supporting an inference is composed of the propositions that were

⁴From the Latin *ponere*, "to affirm".

⁵From the Latin *tollere*, "to deny".

used as the antecedents of the rule that resulted in the inference. If one or more of the evidenciary propositions changes, then the inferred consequent must be examined to verify that it is still true. If an inference is no longer supported by any evidence, then the inference must be denied.

More formally, Doyle restricted the application domain to those systems in which all the propositions which satisfy modus ponens *also* satisfy

$$\begin{array}{l} p \supset q \\ \sim p \\ \therefore \sim q \end{array} \quad (17)$$

Propositions which meet this criterion are said to be logically equivalent. An example of logical equivalence is the state of a switch and the state of the voltage on a line controlled by the switch. When the switch is on, the voltage on the line is on, and while the switch is off, the voltage is off. While there are many examples of natural phenomena which exhibit this kind of behavior, there are many more which do not.

In Doyle's version of truth maintenance, evidence is kept in a "support list". In essence, a support list associated with an inference contains those propositions and a reference to the rule that were used to conclude it. Consider two rules which conclude the same proposition

$$R_1 : p \wedge q \supset c \quad (18)$$

and

$$R_2 : a \wedge b \supset c \quad (19)$$

These two rules are equivalent to a disjunction which embodies both rules

$$R_3 : (p \wedge q) \vee (a \wedge b) \supset c \quad (20)$$

If c is declared true, then its support list will contain one or the other (or both) of the conjunctions in the antecedent. If, at some future time, one of the propositions in the support list changes to false, the conjunction it appears within is removed from the support list. If the support list is empty, then by (17) c must be denied.

This scheme works well for the restricted set of domains which satisfy (17). Unfortunately, most processes occurring in nature **do not** fall into this set, and truth maintenance cannot be used successfully to reason about their dynamics.⁶ Further, there is no explicit mention of time in the truth maintenance mechanism, so the length of time that a proposition might be true cannot be easily specified.

⁶In classical logic, the form described in (17) is considered to be an error, and is known as "denying the premise".

4 Default Logic

Classical logic treats knowledge in a very restricted way. There is a tacit assumption that the reasoner knows all that is necessary in order to proceed. While this may be true in mathematical proof, it is certainly *not* true in general. Human reasoners are often faced with *lack of knowledge*.

There is an ambiguity in the standard meaning of $\sim p$. This can be easily demonstrated. Suppose that p represents the statement, "It is raining". One interpretation of $\sim p$ is the statement, "It is *not* raining". Another interpretation is, "No information is available as to whether it is raining or not". In *either* case, it cannot be said that p is TRUE. For the purpose of mathematical proof, *either* meaning will suffice.

The domain of human reason is not as restricted as that of mathematical logic. The difference between knowing something is FALSE or TRUE and **not knowing** may be crucial. It is very important that the explicit meaning of $\sim p$ be well defined and understood.

Actually, there is implicit to the statement that some proposition is true the further statement that it is *known* to be true. When it is stated that, "It is raining", the actual *meaning* is that, "It is *guaranteed* that it is raining". Similarly, when some proposition is reported false, the meaning is that it is *certainly* NOT true.

Let us look more closely at the meaning of negation. There are statements such as, "Day is *not* night", in which the negation is a property of the domain. That is, in the case that the meaning of negation is the opposite state of the proposition, the veracity of the logical connection is semantically derived.

Alternatively, denial of possession of knowledge concerning the truth of a proposition has nothing to do with the domain of discourse. It is a purely *logical* matter, having only to do with form. Having or not having knowledge about a proposition has nothing to do with the semantic content of the proposition.

The importance of this distinction may be exhibited by reference to a rule used in the Modus Tolens progression (equation 14). If $\sim q$ means that q is *known* to be false, then the result of the progression is that p is *known* false as well. If, instead, $\sim q$ indicates that no knowledge is available about whether q is true (or false), then the result is that p is unknown, too. While the results look formally the same, the meaning is very different. In the case that the state of a proposition is known, we may not logically make any assumption concerning it. But when we have no direct knowledge about a proposition, it is often useful (or indeed, necessary) to make *assumptions*. For example, suppose that it is *known* that the temperature outdoors is 40°F, the dew point is just one degree less, and the humidity high. While it cannot be stated with certainty that there is fog, there is evidence to support the *assumption* that there is fog. Of course, if, in addition to the above, there is specific information that no fog is observable, an assumption about fog

should not be made.

This ability to make assumptions is the utility of supporting an “excluded middle”. As a prerequisite, the two ambiguous meanings of **not** must be formally distinguished. For the remainder of this discussion, “ \sim ” will be used to denote that the opposite of the proposition following it is true, while “ $\bar{\circ}$ ” will signify the lack of verifiable or trustworthy knowledge of the proposition which follows it.

A few rules concerning assumptions are necessary:

- Observation must always take precedence over assumption. This means that if an assumption has been inferred about some proposition, and contradictory information is subsequently observed concerning the same proposition, the assumption must be replaced by the observation.
- All propositions which are concluded from rules in which one or more of the antecedent propositions are assumptions, are themselves assumptions.

A new class of rules called **assumptive** rules may be defined. They are separated from normal rules because all their antecedents contain one (and only one) proposition claiming lack of knowledge about something, and they conclude an assumption. In the case of the “fog” example, let Fx indicate the proposition that there is fog at airport x . Let Qx indicate that the dew point is within two degrees of the outside air temperature at airport x . The assumptive (or default) rule may be written

$$(\bar{\circ} Fx) \wedge Qx \supset Fx^* \quad (21)$$

to mean, “If it is not known that there is fog at an airport, but it is known that the dew point is close to the air temperature at that airport, it can be assumed that there is fog at that airport”. The superscript asterisk appended to the conclusion is a reminder that this proposition is an assumption.

This particular choice of mechanism for making assumptions is a special case of Reiter’s “Default Logic”[4].

The existence of assumptions extends the meaning of contradiction. To appreciate this, it is necessary to understand the different classes of information which can be present in a defeasible logic with default rules:

Observations are propositions that are obtained from outside of the logical system.

Inferred Facts are the consequent propositions resulting from rules in which all antecedent propositions are either observations or inferred facts.

Fundamental Assumptions are the consequent propositions arising from default rules.

Inferred Assumptions are consequent propositions which derive from rules in which one or more antecedent propositions are assumptions (of either variety).

The meaning of contradiction depends on the classes of the propositions involved in the inconsistency:

- Observation and observation: One of the observations must be in error. Given enough domain information (in the form of rules), it may be possible to ascertain which proposition to believe, but this is generally difficult even for humans.
- Observation and Inferred Fact: This situation almost always indicates an erroneous rule. Some rule in the deductive chain leading to the inferred fact must be responsible; with luck, there might only be one.
- Two Inferred Facts: A contradiction of this kind is also indicative of an error in a rule, but the reasoner will require outside assistance to determine the culprit, since there is no way to tell which rule is mistaken.
- Observation *or* Inferred Fact and Fundamental Assumption: The default rule that asserted the fundamental assumption is incorrect.
- Observation *or* Inferred Fact and Inferred Assumption: This is an interesting case. If only one of the propositions in the antecedent of the rule which inferred the assumption is an assumption, then the inferred assumption must be denied, and the assumption which appeared in the antecedent must be denied, and so on back to the fundamental assumption which started that chain of reasoning. This process is called “dependency directed backtracking”. If it should happen that there were more than one assumption in the antecedent of any of the rules in the chain that led to the discrepancy, then a choice must be made: Which assumption should be retracted? There are various approaches that might be taken to answer this question:
 - One could retract the chronologically latest assumption and search for an alternative.
 - One could withdraw the chronologically earliest and search for an alternative.
 - Alternative assumptions could be found for each of the candidates, trying each one until a choice is found which does not cause the contradiction. This may be very time consuming, or, in fact, undecidable.
- Two Fundamental Assumptions: One of the default rules is incorrect.

In essence, all of these possibilities reduce to two major situations.

1. Contradictions among facts, which indicate errors of some kind.
2. Contradictions among assumptions, which require the replacement of one of the propositions with another assumption. The difficulty lies in deciding which assumption to replace, and with what.

The essence of the solution to the problem of replacing an assumption involved in a contradiction is the employment of a class of domain dependent preference rules. These rules have a general form of “If there is a contradiction involving two assumptions regarding X , then Y is the preferred assumption to retain (or to retract).”

The idea of preference rules may also be used to decide which of two (or more) conflicting observations to keep. If each observation is tagged with a description of its origin, then preference rules stating that one origin is more “believable” or more “important” may be used.

5 Temporally Dependent Propositions

Time dependence can be formally introduced into logic by defining

$$p(\tau) \tag{22}$$

to represent that proposition p is true during the time interval τ . The interval τ is a pair of numbers, such as Universal Times, such that the first member of the pair precedes or is equal to (i.e., is *before* or *at the same time of*) the second. Formally,

$$\tau = (t_1, t_2), \quad t_1 \preceq t_2 \tag{23}$$

Having introduced this notion of the interval of veracity or the *activity interval* of a proposition, its effect upon all of the axioms of logic introduced in the previous section must be explored. It will suffice to examine only Conjunction, Disjunction, and the activity interval of the consequent of a rule.

The activity interval of a conjunct will be defined as the *intersection* of the activity intervals of the operands:

$$p(\tau_1) \wedge q(\tau_2) = p \wedge q(\tau_1 \cap \tau_2) \tag{24}$$

That this is a reasonable definition can be seen in the following example: If I am in room A during the time interval from two until four this afternoon ($p(\tau_1)$), and you are in room A during the interval from three until five this afternoon ($q(\tau_2)$), then *we* are in room A from three until four this afternoon ($p \wedge q(\tau_1 \cap \tau_2)$).

Nota Bene: This definition of conjunction effectively states that two propositions can interact if and only if they have overlapping time intervals. This may seem overly restrictive at first glance, especially considering the human penchant for describing many interacting events as following one another and being causally linked. However, closer examination reveals that the restriction is completely correct. Temporally disjoint events which *seem* to interact are invariably connected by some persistent process, produced as

an effect of the first event, which remains in effect at least until its time interval overlaps that of the second event.

The preceding motivates the definition of the activity interval of a disjunctive pair as the *union* of the activity intervals of the operands:

$$p(\tau_1) \vee q(\tau_2) = p \vee q(\tau_1 \cup \tau_2) \quad (25)$$

Again, using the room occupancy example: If I am in room *A* during the time interval from two until four this afternoon ($p(\tau_1)$), and you are in room *A* during the interval from three until five this afternoon ($q(\tau_2)$), then one or the other of us is in room *A* from two until five this afternoon ($p \vee q(\tau_1 \cup \tau_2)$).

5.1 Evanescence and Persistence

The rules for generating the activity interval of a logical combination of temporally constrained propositions allow the computation of the activity interval of the antecedent of a rule. However, the activity interval of the consequent of a rule is not necessarily the same as that of its antecedent. Processes and physical things whose activity intervals are shorter than the activity intervals of their antecedents are called *evanescent*. An example of an evanescent process is the firing of a gun. When the hammer drops the gun fires. The fact that the hammer remains down does not make the “bang” last longer. Other things and processes are persistent, lasting well after the events which created them have ceased to exist. For example, if the temperature is below freezing on the ground, and it is raining, ice will form on the ground. When the rain stops, the ice does not simply disappear.

The root of this problem is the domain dependence of the activity interval of a physical process. While it is true that no process can exist without some form of causal precedent, once formed a process may have an independent existence of its own. In the case of persistence, quite often the only way to “undo” something which has been “done” is to do something specifically designed to destroy it.

For practical purposes, there are only two classes of consequent

- Inferred propositions with activity intervals which are the same as the computed activity intervals of their antecedents. In this case, no further information is required in the consequent.
- Inferred propositions with independent activity intervals. Causality requires that the beginning of the consequent activity interval be the same as the start of the computed antecedent interval, but the domain dependent information to compute the end of the consequent activity interval must be supplied in the rule.

This topic will be discussed in greater detail in a forthcoming paper, “Implementation of the Tower Chief Planning System”.

5.2 Resource Allocation and Planning

Classical temporal logic is insufficient to describe the domain and events which occur in the ATC environment characterized previously. In particular, the information available about the future state of some value may change during the execution of a plan. For example, there may be a weather prediction at 09:00 that claims that passage of a front with an associated shift in wind direction will occur between noon and one o'clock. A later prognostication, perhaps at 11:00, might change the time of the frontal passage or some parameter associated with it. Such a change may require a modification of a planned configuration shift which may already be in progress. Because of the infeasibility of certain transitions, the modification of one configuration choice may affect those which precede and follow it, and so on.

Preparation for the use of a specific configuration may demand the allocation of resources in advance. In winter, for example, one or more of the runways that are to be used in a future configuration may require snow or ice removal or treatment to prevent ice accumulation prior to being put into service.

In standard expositions on temporal logic, the processes that the system is designed to model usually involve the evolution of some physical quantity such as the position of a ball or the temperature of an object. The rules for this kind of modelling generally look like

$$p(t_1, t_2) \wedge q(t_3, t_4) \supset r(t_5, t_6) \quad (26)$$

where $p(t_1, t_2) \wedge q(t_3, t_4) \equiv p \wedge q [(t_1, t_2) \cap (t_3, t_4)]$. If r is a new state that was described previously by p , then $t_2 \preceq t_5$.⁷ For example, "If some water is in the liquid state (p) during (t_1, t_2) , and is brought in contact with a thermally massive object with a temperature greater than the boiling point of water (q) during (t_3, t_4) , then the water will be in the vapor phase (r) during (t_5, t_6) ."

A plan, on the other hand, is by definition something that is intended for future execution. Plans are based on what the planner *believes* is going to happen. Because the future is not fixed, a plan may have to be modified or even abandoned before or during its execution. Actions which are taken for the most part cannot be withdrawn when a plan is deserted. That such actions were based on beliefs which turned out to be false does not change the fact that they were executed. If the planner is to be able to explain the reasons behind its actions, it must recall its prior beliefs even if they were later proven to be wrong. In fact, the planner may change what it believes about the future as a result of its intention to carry out some plan.

Consider the reasoning involved in the allocation of some resource. Let $\mathcal{W}r(t_1, t_2)$ indicate that use of resource r is desired during the time interval (t_1, t_2) . Similarly, let $\mathcal{A}r(t_1, t_2)$ indicate that resource r is available during the time interval, and finally let

⁷In other words, a thing cannot be in two disjoint states at the same time.

$\mathcal{H}r(t_1, t_2)$ represent that the resource has been allocated for the period. Then we might write

$$\mathcal{W}r(t_1, t_2) \wedge \mathcal{A}r(t_1, t_2) \supset \mathcal{H}r(t_1, t_2) \wedge \sim \mathcal{A}r(t_1, t_2) \quad (27)$$

to describe the rule for allocation: “If resource r is required during (t_1, t_2) , and the resource is available during the period, then r is allocated for the interval, and is no longer available for allocation during that time.” There is an apparent paradox in this formulation, since $\mathcal{A}r(t_1, t_2)$ and $\sim \mathcal{A}r(t_1, t_2)$ appear at the same time. The reason that this problem appears is that the reasoner’s belief about the availability of the resource for other use during the specified interval changed as a consequence of the reasoning process itself.

The last example demonstrates that reasoning about the future may involve non-monotonic logic. However, propositions previously thought to be true cannot necessarily be simply retracted. They may be in the support lists of later propositions which involve actions or the reasoning process itself. A human reasoner in the resource allocation example would explain, “Of course, I believed the resource was available *before* I allocated its use. Now that I have committed the use of it during the time period, it is no longer available for other use.” The reasoner is aware of the temporal order of the events of the reasoning process in addition to the projected time of execution of the plan.

There are *three* time periods associated with the resource allocation problem. First, the time in the future during which the use of the resource is desired. For the sake of easy reference, let us give that a name: the *activity interval*. Second, the time during which the planner believes that the resource will be available during the activity period. Third, the time during which the planner believes that the resource has been allocated for the activity period and will no longer be available for other use during that interval. Again, for reference, let’s call these two periods “belief intervals”. If we put the belief intervals into the logical statement of the resource allocation rule as subscripts to the propositions, we obtain

$$\mathcal{W}r_{(\tau_1, \tau_2)}(t_1, t_2) \wedge \mathcal{A}r_{(\tau_3, \tau_4)}(t_1, t_2) \supset \mathcal{H}r_{(\tau_5, \tau_6)}(t_1, t_2) \wedge \sim \mathcal{A}r_{(\tau_5, \tau_6)}(t_1, t_2) \quad (28)$$

and the paradox is resolved. The planner can now refer to what it believed prior to making its decision as well as its opinion after the act of making the decision.

Since a runway manager’s duties include allocation of people and equipment to a variety of tasks, and a manager must take *actions* based on the current knowledge about the domain, the kind of reasoning described above is central to accomplishing the cognitive task of planning runway configurations.

6 Conclusions

The important attributes of this problem are not unique to runway configuration management. Temporal reasoning, default reasoning, and reasoning about the commitment

of resources are ubiquitous characteristics of almost all process management. Analysis of these kinds of cognition led us to a single representation and reasoning paradigm which integrates all three.

There is still much to be done. Currently, all input activity intervals must be clock times. It would be much more convenient to be able to enter times in qualitative terms by reference to information already known. Clock times often give a false sense of precision to information whose actual accuracy is fuzzy at best. It is well known that the process of maintaining a temporal database of the kind described here is \mathcal{NP} -hard.⁸ What saves us is that the rate of change in the systems we have looked at is slow, and the airport effectively hardware resets every night. If this technology is to be employed in a broader spectrum of application, such as planning and scheduling of a planetary exploration robot, this efficiency problem will have to be solved.

7 Acknowledgements

I would like to thank the staff and administration of the Boston TRACON for their invaluable cooperation and assistance. Thanks also to the staff of the Flight Transportation Laboratory, especially Dennis F. X. Mathaisel and John D. Pararas, for many beneficial discussions.

⁸ \mathcal{NP} -hard refers to the computational complexity of the problem. \mathcal{NP} means “Non-deterministic Polynomial, indicating that the computation time for the brute force approach to the computation rises exponentially with the size of the problem. See [10].

References

- [1] I. M. Bocheński. *A History of Formal Logic*. Chelsea Publishing Co., New York, 1970.
- [2] Irving M. Copi. *Introduction to Logic*. MacMillan, New York, 1972.
- [3] Jon Doyle. A Truth Maintenance System. *Artificial Intelligence*, 12:231–272, 1979.
- [4] Raymond Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.
- [5] James F. Allen. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832–843, November 1983.
- [6] James F. Allen and J. A. Kooman. Planning using a temporal world model. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 741–747, 1983.
- [7] T. Dean and K. Kanazawa. Probabilistic causal reasoning. In *Proceedings of the Fourth Workshop on Uncertainty in Artificial Intelligence*, pages 73–80, 1988.
- [8] Drew McDermott. A temporal logic for reasoning about processes and plans. *Cognitive Science*, 6:101–155, 1982.
- [9] Yoav Shoham. *Reasoning about Change: Time and Causation from the Standpoint of Artificial Intelligence*. MIT Press, 1988.
- [10] G. Edward Barton, Robert C. Berwick, and Eric Sven Ristad. *Computational Complexity and Natural Language*. MIT Press, Cambridge, Massachusetts, 1988.